

10/28/99

JCS05 U.S. PTO

Please type a plus sign (+) inside this box ☐

A

UTILITY PATENT APPLICATION TRANSMITTAL (Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))		Attorney Docket No. 15311-2207
First Inventor or Application Identifier		Stephen Dickson
Title		COMPUTERIZED FILE SYSTEM AND METHOD
Express Mail Label No.		EL024695925US

APPLICATION ELEMENTS See MPEP chapter 600 concerning utility application contents	ADDRESS TO: Assistant Commissioner for Patents Box Patent Application Washington, DC 20231
---	--

<p>1. <input checked="" type="checkbox"/> *Fee Transmittal Form (e.g., PTO/SB/17) (Submit an original and a duplicate for fee processing)</p> <p>2. <input checked="" type="checkbox"/> Specification [Total Pages 26] (preferred arrangement set forth below)</p> <ul style="list-style-type: none">- Descriptive title of the invention- Cross References to Related Applications- Statement Regarding Fed sponsored R & D- Reference to Microfiche Appendix- Background of the invention- Brief Summary of the invention- Brief Description of the Drawings (if filed)- Detailed Description- Claim(s)- Abstract of the Disclosure <p>3. <input checked="" type="checkbox"/> Drawing(s) [Total Sheets 6]</p> <p>4. Oath or Declaration [Total Pages 1]</p> <p>a. <input checked="" type="checkbox"/> Newly executed (original copy)</p> <p>b. <input type="checkbox"/> Copy from a prior application (37 C.F.R. § 1.63(d)) (for continuation/divisional with Box 17 completed) [Note Box 5 below]</p> <p>DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).</p> <p>5. <input type="checkbox"/> Incorporation By Reference (useable if Box 4b is checked) The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered to be part of the disclosure of the accompanying application and is hereby incorporated by reference therein</p>	<p>6. <input type="checkbox"/> Microfiche Computer Program (Appendix)</p> <p>7. Nucleotide and/or Amino Acid Sequence Sequence Submission (if applicable, all necessary)</p> <p>a. <input type="checkbox"/> Computer Readable Copy</p> <p>b. <input type="checkbox"/> Paper Copy (Identical to computer copy)</p> <p>c. <input type="checkbox"/> Statement verifying identity of above copies</p> <p>ACCOMPANYING APPLICATION PARTS</p> <p>8. <input checked="" type="checkbox"/> Assignment Papers (cover sheet & document(s))</p> <p>9. 37 C.F.R. § 3.73(b) Statement (when there is <input type="checkbox"/> Power of Attorney an assignee)</p> <p>10. <input type="checkbox"/> English Translation Document (if applicable)</p> <p>11. <input type="checkbox"/> Information Disclosure Statement (IDS)/PTO-1449 <input type="checkbox"/> Copies of IDS Citations</p> <p>12. <input type="checkbox"/> Preliminary Amendment</p> <p>13. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503) (Should be specifically itemized)</p> <p>14. <input type="checkbox"/> *Small Entity Statement(s) <input type="checkbox"/> Statement filed in prior application, Status still proper and desired (PTO/SB/09-12)</p> <p>15. <input type="checkbox"/> Certified Copy of Priority Document(s) (if foreign priority is claimed)</p> <p>16. <input type="checkbox"/> Other:</p>
--	---

*NOTE FOR ITEMS 1 & 14: IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C.F.R. § 1.27), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.28).


17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information below and in a preliminary amendment:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: /

Prior application Information: Examiner

Group/Art Unit:

18. CORRESPONDENCE ADDRESS					
<input type="checkbox"/> Customer Number or Bar Code Label			or <input checked="" type="checkbox"/> Correspondence address below		
(Insert Customer No. or Attach bar code label here)					
Name	Christopher K. Gagne				
Address	Cesari and McKenna 30 Rows Wharf				
City	Boston	State	MA	Zip Code	02110
Country	U. S. A.	Telephone	(617) 951-2500	Fax	(617) 951-3927

Name (Print/Type)	Christopher K. Gagne	Registration No. (Attorney/Agent)	36,142
Signature		Date	October 28, 1999

JCS05 U.S. PTO
09/428364

10/28/99

In Re The Application of:
Stephen Dickson
Serial No.: Not yet assigned
Filed: October 28, 1999
For: COMPUTERIZED FILE SYSTEM
AND METHOD

Art Unit: Not Yet Assigned

EXPRESS-MAIL DEPOSIT

The following Patent Application (26 pages including 36 claims), Formal Drawings (6 sheets), Utility Patent Application Transmittal Letter, Fee Transmittal Letter, Assignment, Recordation Form Cover Sheet, Declaration for Patent Application, and a check in the amount of \$1,712 are being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service pursuant to 37 C.F.R. §1.10, in an envelope addressed to Box Patent Application, Assistant Commissioner for Patents, Washington, D.C. 20231, on October 28, 1999.

UNITED STATES PATENT APPLICATION

of

Stephen Willard Dickson

for

COMPUTERIZED FILE SYSTEM AND METHOD

COMPUTERIZED FILE SYSTEM AND METHOD

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention relates to a computer file system and method, wherein one
5 or more characteristics (e.g., actual data contained in the file, and/or meta-data associated
with the file, such as file name/handle, ownership, related links, size, time of last modifi-
cation, user access privilege-related data, etc.) of a data file maintained by a computer
process (e.g., residing in a network computer node) may be accessed and modified by
multiple other computer processes (e.g., residing in computer network client nodes), and
10 a mechanism exists to maintain the coherency of the data file and its characteristics de-
spite their being subject to modification by the multiple processes in the network. As
used herein, the term "data file" comprises objects in a distributed computer file system,
such as user and system program and data files, directories, and associated objects. Also
as used herein, the "modification" of a data file may comprise the creation of the data
15 file.

Brief Description of Related Prior Art

Data communication in a computer network involves data exchange between two
or more entities interconnected by communication links. These entities are typically
software program processes executing on computer nodes, such as endstations and inter-
20 mediate stations. Examples of an intermediate station may be a router or switch which

interconnects the communication links and subnetworks to enable transmission of data between the endstations. A local area network (LAN) is an example of a subnetwork that provides relatively short distance communication among the interconnected stations, whereas a wide area network enables long distance communication over links provided
5 by public or private telecommunications facilities.

Communication software executing on the endstations correlate and manage data communication with other endstations. The stations typically communicate by exchanging discrete packets or frames of data according to predefined protocols. In this context, a protocol consists of a set of rules defining how the stations interact with each other. In
10 addition, network routing software executing on the routers allow expansion of communication to other endstations. Collectively, these hardware and software components comprise a communications network and their interconnections are defined by an underlying architecture.

Modern communications network architectures are typically organized as a series
15 of hardware and software levels or "layers" within each station. These layers interact to format data for transfer between, e.g., a source station and a destination station communicating over the network. Predetermined services are performed on the data as it passes through each layer and the layers communicate with each other by means of the predefined protocols. The lower layers of these architectures are generally standardized and
20 are typically implemented in hardware and firmware, whereas the higher layers are generally implemented in the form of software running on the stations attached to the network. In one example of such a communications architecture there are five layers which are termed, in ascending interfacing order, physical interface, data link, network, trans-

port and application layers. These layers are arranged to form a *protocol stack* in each communicating station of the network. Figure 1 illustrates a schematic block diagram of prior art protocol stacks 125 and 175 used to transmit data between a source station 110 and a destination station 150, respectively, of a network 100. As can be seen, the stacks
5 125 and 175 are physically connected through a communications channel 180 at the interface layers 120 and 160. For ease of description, the protocol stack 125 will be described.

In general, the lower layers of the communications stack provide internetworking services and the upper layers, which are the users of these services, collectively provide
10 common network application services. The application layer 112 provides services suitable for the different types of applications using the network, while the lower interface layer 120 accepts industry standards defining a flexible network architecture oriented to the implementation of LANs.

Specifically, the interface layer 120 comprises the physical interface layer 126,
15 which is concerned with the actual transmission of signals across the communication channel and defines the types of cabling, plugs and connectors used in connection with the channel. The data link layer (i.e., "layer 2") 121 is responsible for transmission of data from one station to another and may be further divided into two sublayers: Logical Link Control (LLC 122) and Media Access Control (MAC 124).

20 The MAC sublayer 124 is primarily concerned with controlling access to the transmission medium in an orderly manner and, to that end, defines procedures by which the stations must abide in order to share the medium. In order for multiple stations to share the same medium and still uniquely identify each other, the MAC sublayer defines

a hardware or data link address called a MAC address. This MAC address is unique for each station interfacing to a LAN. The LLC sublayer 122 manages communications between devices over a single link of the network.

The network layer 116 (i.e., "layer 3") provides network routing and that relies on transport protocols for end-to-end reliability. An example of a network layer protocol is the Internet Protocol ("IP"). An example of such a transport protocol is the Transmission Control Protocol (TCP) contained within the transport layer 114. The term *TCP/IP* is commonly used to refer to the Internet architecture. (See, e.g., Tanenbaum, Computer Networks, Third Ed., Prentice Hall PTR, Upper Saddle, River, New Jersey, 1996).

Data transmission over the network 100 therefore consists of generating data in, e.g., sending process 104 executing on the source station 110, passing that data to the application layer 112 and down through the layers of the protocol stack 125, where the data are sequentially formatted as a frame for delivery onto the channel 180 as bits. Those frame bits are then transmitted over an established connection of channel 180 to the protocol stack 175 of the destination station 150 where they are passed up that stack to a receiving process 174. Data flow is schematically illustrated by solid arrows.

Although actual data transmission occurs vertically through the stacks, each layer is programmed as though such transmission were horizontal. That is, each layer in the source station 110 is programmed to transmit data to its corresponding layer in the destination station 150, as schematically shown by dotted arrows. To achieve this effect, each layer of the protocol stack 125 in the source station 110 typically adds information (in the form of a header) to the data generated by the sending process as the data descends the stack.

For example, the network layer encapsulates data presented to it by the transport layer within a packet having a network layer header. The network layer header contains, among other information, source and destination (logical) network addresses needed to complete the data transfer. The data link layer, in turn, encapsulates the packet in a frame
5 that includes a data link layer header containing information required to complete the data link functions, such as (physical) MAC addresses. At the destination station 150, these encapsulated headers are stripped off one-by-one as the frame propagates up the layers of the stack 175 until it arrives at the receiving process.

A computer file system controls the formatting of data files, maintaining the loca-
10 tion of the data files in memory, the logical hierarchy of data files, user/process access privileges (e.g., in terms of reading and writing) to the data files, and other file-related tasks, such as house-keeping and administrative functions that keep track of data file statistics (e.g., sizes of the files, dates of creation and last modification of the files, etc.). Computer file systems are frequently integrated with the operating system such that, al-
15 though a logical or functional distinction may be made between the two systems, they are intertwined with each other from a source code standpoint. When the processes that implement the file system reside in multiple nodes in a computer network, that file system may be termed a "distributed" computer file system.

A "client/server network" is one conventional type of computer network archi-
20 tecture wherein data files stored or residing in one computer node (commonly termed a "server" computer) in the network are shared, using a distributed computer file system, by multiple processing executing/residing in other computer nodes (commonly terms "client" computers) in the network. That is, data files and their characteristics stored or

residing in the server computer node may be accessed and modified, via the distributed file system, by multiple processes executing/residing in the client computer needs.

5 The client/server network architecture offers advantages over other types of network architectures. For example, since in a client/server network, data files residing in the server computer node may be accessed by processes residing in the client computer nodes, copies of these files need not also reside in the client nodes. This increases the amount of client computers' resources that may be made available for other purposes, and eliminates the cost and time necessary to support and maintain separate copies of these files in the client computers.

10 In distributed file systems, maintaining the coherency of data files and file characteristics shared among, and subject to modification by multiple processes residing in the client nodes can be problematic. That is, since multiple processes residing in the client nodes may be able to access and modify the characteristics of data files stored in the server node, it becomes necessary for the file system to ensure coherency of these characteristics despite their being subject to modification by the multiple processes.

15 In one conventional solution to this problem, a file system management process residing in the server node grants sets (i.e., combinations) of different types of "tokens" to requesting client node processes that grant permission to the processes to modify particular characteristics of files stored in the server node. Each "token" is identified by the particular class/type to which it belongs, and is associated with a respective data file. In order for a process to be able to execute a respective modification to a respective data file characteristic, the process must first be granted permission by the network server's file

management process, in the form of a grant of a respective set of different types/classes of tokens associated with that file and the modification.

More specifically, according to this prior art technique, when a client node process desires to modify a respective characteristic of a respective data file stored in the server node, the process transmits separate respective token grant request messages to the server node's file management process that request grant of each of tokens in the predetermined set of tokens required for permission to make the desired modification. In response to each respective request message, the file management process determines whether the respective token whose grant is being requested by the respective request message is available for grant to the client node process. If the respective token is available for grant, the file management process transmits a token grant message to the client process that grants that respective token to the client process. Conversely, if the respective token is not available for grant, for example, as a result of being currently granted to another client node process, the file management process may transmit a token revocation message to the other client node process to which the respective token is currently granted. In response to the token revocation message, the other client node process forwards to the file management process a token relinquishment message indicating that the other client node process has relinquished its grant of the respective token, thereby returning the respective token to the pool of tokens available for grant to the requesting client node process. The file management process may then transmit the token grant message to the client process. A client node process may execute a desired modification to a respective data file only after, and for as long as, the process has been granted the respective set of tokens required to make the desired modification.

Although this prior art technique is able to maintain the coherency of data file characteristics stored in the server node, it has certain disadvantages and drawbacks. For example, since only a single respective token may be requested and granted in each token request and grant message, respectively, when the set of tokens required for a desired file characteristic modification comprises more than one token, multiple token request and grant messages must be exchanged between the file management and requesting client node processes to enable the requesting client node process to carry out the desired file characteristic modification. Likewise, since only a single respective token may be revoked and relinquished in each token revocation and relinquishment message, respectively, if multiple tokens must be revoked and relinquished to enable the desired modification to take place, multiple token revocation and relinquishment messages must be exchanged to effect the revocation and relinquishment of such multiple tokens. Thus, since, at any given time, many client node processes may seek to modify, and may presently be engaged in modification of, characteristics of frequently-used data files stored in the server node, this can result in an undesirably large amount of network bandwidth being consumed by tasks related to network file system overhead, and can undesirably increase network congestion.

SUMMARY OF THE INVENTION

In accordance with the present invention, a computerized file system and method are provided that are able to overcome the aforesaid and other disadvantages of the prior art. More specifically, in one embodiment of the system of the present invention, a distributed computer file system is provided for use in a client/server network. A first process (e.g., a file system management process) residing in a server computer node maintains

a data file in computer-readable memory in the server computer node. A second process (e.g., a client node process) generates a first message that requests that the second process be granted by the first process a plurality of tokens required for the second process to be able to modify at least one characteristic of the data file. In response to the first message,
5 the first process generates a second message that grants the tokens to the first process if the tokens are available for grant to the second process.

In this embodiment, if any of the tokens are unavailable for grant to the second process as a result of being currently granted to a third process (e.g., a process residing in a second computer node), the first process may generate a third message that requests that
10 the third process relinquish its current grant of the unavailable tokens. In response to the third message, the third process may generate a fourth message relinquishing the current grant of these tokens, thereby making the tokens available for grant by the first process to the second process. The second and third processes may reside in two different client nodes that are geographically remote from each other.

In accordance with the present invention, the tokens whose grant are requested by the second process via the first message, and granted to the second process via the second message, respectively, may comprise all of the tokens required to enable the second process to be able to modify the at least one characteristic of the data file. Further, the third and fourth messages each may specify a plurality of tokens whose previous grant is to be
15 revoked and relinquished, respectively. This permits the number of token request, grant, revocation, and relinquishment messages that need to be generated and transmitted according to the present invention to be substantially reduced compared to the prior art. Advantageously, this permits the amount of network bandwidth that is consumed by tasks
20

related to network file system overhead in a network embodying the present invention to be substantially reduced compared to the prior art. Further advantageously, this may also permit the amount of network congestion present in a network embodying the present invention to be reduced compared to the prior art.

5 These and other advantages of the present invention will become apparent as the following Detailed Description proceeds and upon reference to the Drawings, wherein like numerals depict like parts, and in which:

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a highly schematic block diagram of prior art communications archi-
10 tecture protocol stacks used to transmit data between stations of a computer network.

Figure 2 is a highly schematic block diagram of a computer network comprising one embodiment of the computer file system of the present invention.

Figure 3 is a diagram illustrating messages that may be transmitted and received by processes in the network of Figure 2.

15 Figure 4 is a highly schematic diagram illustrating processes and related program instructions and data structures that reside in memory in the server node in the network of Figure 2.

Figure 5 is a highly schematic diagram illustrating processes and related program instructions and data structures that reside in memory of one of the client nodes in the
20 network of Figure 2.

Figure 6 is a highly schematic diagram illustrating processes and related program instructions and data structures that reside in memory of another one of the client nodes in the network of Figure 2.

Although the following Detailed Description will proceed with reference being
5 made to illustrative embodiments and methods of use, it should be understood that the present invention is not intended to be limited to these embodiments and methods of use. Rather, the present invention should be viewed broadly, as being of broad scope limited only as set forth in the hereinafter appended claims.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENT

10 With reference being made to Figures 2-6, a distributed computer network 200, wherein one embodiment of the present invention may be advantageously used, will now be described. Network 200 includes a server computer node 202 and a plurality of client computer nodes 204, 206 interconnected by a network communications system (represented by network cloud 208). In general, each of the computer nodes includes com-
15 puter-readable memory 210 for storing software programs, algorithms, and data structures associated with, and for carrying out, the inventive file system, and related and other methods and techniques described herein. In addition, each of the nodes further includes processor 212 for executing these software programs and algorithms, and for manipulating the stored data structures, to enable the nodes to carry out these methods
20 and techniques in network 200. Each of the client nodes 204, 206 may also include a conventional user input/output interface 214 (e.g., comprising keyboard, pointing device, display terminal, etc.) for permitting a user to control and interact with the node. Like-

wise, although not shown in the Figures, the server node 202 may also comprise such a user interface.

In operation, a data file and its associated characteristics (collectively and/or singly referred to herein after by the term "data file characteristic" and by numeral 250) are stored in memory 210 of node 202. The data file characteristic 250 is maintained in the server node's memory 210 by file system management process and associated executable instructions 252 resident in the server node's memory 210. The process 252 may be spawned by (or, alternatively, comprised in) one or more executing operating system processes 254 (e.g., comprised in the Alpha Server 64 Unix 5.0™ operating system of Compaq Computer Corp. of Houston, Texas) that may also be resident in the memory 210 of the server node 202.

Also in operation, an executing distributed file system process 260 is resident in the memory 210 of client node 204. This process 260 may be spawned by (or, alternatively comprised in) one or more executing operating system processes and associated executable instructions 262 also resident in the memory 210 of the client node 204. Cache memory 264, the purpose of which will be described in greater detail below, may also be comprised in the memory 210 of client node 204.

Also in operation, an executing distributed file system process 280 is resident in the memory 210 of client node 206. This process 280 may be spawned by (or, alternatively comprised in) one or more executing operating system processes and associated executable instructions 282 also resident in the memory 210 of the client node 206. Cache memory 284, the purpose of which will be described in greater detail below, may also be comprised in the memory 210 of client node 206.

In this embodiment of the system of the present invention, when distributed file system process 260 desires to modify the data file characteristic 250 in server 202 (e.g., as a result of user input of commands to process 260 via the user interface of node 204), process 260 first determines which tokens are required to be granted to process 260 in order for process 260 to be able to carry out the desired modification to data file characteristic 250. For example, in system 200, in order to write to file 250 and thereby change its characteristics, process 260 must be granted by process 252 the following tokens: the ATTR_TOKEN token which grants permission to locally cache in cache memory 264 the non-static portions of the meta-data associated with the file 250, the MTIME_TOKEN token that grants permission to locally cache in cache memory 264 the last modification time of the file, and the SIZE_TOKEN token that grants permission to cache in cache memory 264 the byte size of the file 250. Depending upon the particular type of modification to the file characteristic 250 being requested by process 260, the number and type of tokens required to be granted to process 260 varies. In system 200, each type of respective possible modification that can be made by process 260 to the characteristics of file 250 is associated with a respective, set of different types of tokens, and can only be carried out by process 260 when that process 260 is granted the respective set of tokens associated with the respective modification. Process 260 maintains a table (not shown) in memory 210 of node 204 that correlates each of the respective possible modifications that can be made to the file 250, with the respective sets of tokens that must be granted to process 260 in order for process 260 to be able to carry out the respective modifications. For purposes of the present discussion, the tokens that must be granted to process 260 by

process 252 to grant process 260 permission to execute the desired modification to data file characteristic 250 will be symbolically represented hereinafter as tokens A, B, and C.

When process 260 has determined the token set A, B, C that must be granted to process 260 in order for process 260 to be able to execute the desired modification to file characteristic 250, process 260 generates and forwards a message 300 via network 208 to file system management process 252 in node 202. The message 300 requests that process 260 be granted all of the tokens A, B, C necessary for process 260 to be able to execute the file characteristic modification that it desires to make. Message 300 contains information specifying the tokens, A, B, C, the node 204 and/or process 260 generating the message 300, and the file 250 associated with the tokens A, B, C. Of course, as will be appreciated by those skilled in the art, depending upon the manner in which system 200 is implemented, message 300 may contain information in addition to the foregoing.

Management process 252 maintains a table 251 that correlates respective pools of tokens that can be permissibly granted for each respective data file characteristic modification in the file system with the respective data files with which they are associated, and also indicates which of the tokens is presently granted, and if granted, the process (e.g., 260 or 280) in system 200 to which the token or tokens are granted. When server node 202 receives message 300, process 252 determines from this table 251 which of the tokens requested by the message 300 is currently available for grant to process 260. If as, is shown in the case 1 example in Fig. 3, all of tokens A, B, and C are presently available for grant to process 260 when process 252 receives message 300, process 252 generates and transmits, via network 208, token grant message 302 to process 260 in node 204. Grant message 302 identifies the data file 250 desired to be modified, and the tokens A,

B, C being granted to process 260. Process 252 then updates the information contained in the table 251 from which it determined whether tokens A, B, C were available for grant, to reflect that tokens A, B, and C associated with file 250 have been granted to process 260.

5 Process 260 maintains in cache memory 264 a table 259 that correlates the tokens that it have been granted by process 252 with the data files with which they are associated. When process 260 receives message 302, process 260 updates this table 259 to reflect that process 260 has been granted the tokens A, B, C identified in message 302, and that tokens A, B, C are associated with file 250. Process 260 then determines whether it
10 has yet to be granted any of the tokens (i.e., A, B, or C) that it requires to be able to carry out its desired file characteristic modification, based upon the information contained in the table 259. If after receipt by process 260 of message 302, as is the case in the present example, no additional tokens need to be granted for process 260 to be able to carry out its desired file characteristic modification, process 260 then undertakes such modifica-
15 tion. Conversely, if any such required token has yet to be granted, process 260 waits to make such modification until it has been granted all such required tokens.

Also conversely, as is shown in the case 3 example in Fig. 3, if one or more of the tokens A, B, and/or C whose grant is requested in message 300 is not available for grant to process 260 when process 252 receives message 300 as a result of being currently
20 granted to another client node process (e.g., process 280 in node 206), the server node process 252 generates and transmits to the other client node process 280 a token revocation message (e.g., message 314 in the case 3 example in Fig. 3). Message 314 requests that the other client node process 280 relinquish its grant of those tokens A, B, C whose

grant has been requested by message 300, but currently are granted to process 280.

When process 280 receives the token revocation message 314, the process 280 completes any on-going modification to the data file characteristic 250 that it is presently engaged in that requires grant of any of the tokens A, B, C whose relinquishment has been requested
5 by message 314. Thereafter, process 280 generates and transmits to process 252, via network 208, a token relinquishment message (e.g., message 316 in the case 3 example in Fig. 3) that relinquishes grant of the tokens A, B, C whose relinquishment was requested in message 314. Process 280 then updates table 285 that it maintains in memory 284 that correlates the tokens that process 280 has been granted with the data files associated with
10 those tokens. More specifically, process 280 updates table 285 to reflect the relinquishment of the grant of tokens A, B, C.

When process 252 receives the relinquishment message 216, it generates and transmits to process 260, via network 208, grant message 302, and updates the token grant table 251, in the manner described previously in connection with the case 1 example in Fig. 3. Likewise, when process 260 receives message 302, it updates the token
15 grant table 259, in the manner described previously in connection with the case 1 example in Fig. 3.

Further conversely, as shown in the case 2 example in Fig. 3, if process 260 determines, prior to generating request message 300, that process 260 has already been
20 granted one of the tokens C necessary for it to be able to undertake its desired file characteristic modification, instead of generating and transmitting token request message 300, process 260 generates and transmits to process 252 a different token request message 304. Token request message 304 is the same as request message 300, except that, the

only tokens whose grant is requested by message 304 are those tokens A, B that process 260 determines are not currently granted to process 260. Just as is the case in the previous examples of cases 1 and 3, when process 252 receives message 304, process 252 examines the information in the token grant table 251 that it maintains in memory 210 of node 202 to determine which tokens whose grant is requested in message 304 are presently available for grant to process 260 (e.g., token A in the case 2 example). If process 252 determines that one or more tokens B requested in message 304 is currently unavailable for grant as a result of being currently granted to another client node process (e.g., process 280), instead of generating and transmitting revocation message 314, process 252 may generate and transmit to the other process 280 a different token revocation message 306. The information contained in message 306 is the same as that contained in message 314, except that the only tokens B whose grant message 306 requests be relinquished are those whose grant has been requested in message 304 but are currently granted to process 280.

When process 280 receives message 306, it first completes any on-going modification to the data file characteristic 250 that it may be presently engaged in that requires grant of the token or tokens B whose grant message 306 has requested be relinquished. After completing any such modification to the file characteristic 250, process 280 then generates and transmits to process 252 a token relinquishment message 308 that is the same as the relinquishment message 316, except that relinquishment message 308 only relinquishes grant of the token or tokens B requested by message 306. When process 280 generates and transmits message 308 to process 252, process 280 also updates the token grant table 285 that it maintains in memory 284 to reflect that process 280 no longer is

granted the token or tokens B whose grant has been relinquished by message 308. When process 252 receives message 308, process 252 generates and transmits to process 260 a token grant message 310 granting tokens A, B to process 260. When process 252 generates and transmits message 310 to process 260, process 252 also updates the token grant
5 table 251 to reflect the grant of these tokens A, B to process 260 and the relinquishment of the grant of tokens A, B to process 280.

When process 260 receives message 310, process 260 updates the token grant table 259 to reflect that process 260 has been granted the tokens A, B. Process 260 then undertakes the modification that it desires to make to the data file 250 that requires grant
10 of tokens A, B, and C.

The foregoing description has been directed to specific embodiments of the present invention. It will be apparent, however, that other alternatives, variations and modifications may be made to the described embodiments, with the attainment of some or all of their advantages. Therefore, it is intended that the present invention be viewed as being
15 of broad scope and as covering all such alternatives, modifications, and variations. Thus, it is intended that the present invention be defined only as set forth in the hereinafter appended claims.

20 What is claimed is:

CLAIMS

- 1 1. A computerized data file system, comprising:
2 a first process that maintains a data file in computer-readable memory; and
3 a second process that generates a first message requesting that said second process
4 be granted by said first process a plurality of tokens required for said second process to
5 modify at least one characteristic of said file;
6 said first process generating a second message, in response to said first message,
7 that grants said tokens to said first process if said tokens are available for grant to said
8 second process.
- 1 2. A system according to claim 1, wherein:
2 said first process is resident at a server computer node, and said second process is
3 resident at a client computer node.
- 1 3. A system according to claim 1, wherein:
2 if any of said tokens are unavailable for grant to said second process as a result of
3 current grant of said tokens to at least one other process, said first process generates a
4 third message revoking the current grant of said tokens to said at least one other process.
- 1 4. A system according to claim 3, wherein:
2 said at least one other process, in response to said third message, generates a
3 fourth message making said tokens available for grant by said first process.
- 1 5. A system according to claim 3, wherein:
2 said first process resides in a first computer node;
3 said second process resides in a second computer node;
4 said at least one other process resides in at least one other computer node; and
5 said first computer, second computer, and at least one other computer nodes are
6 networked together and are remote from each other.

1 6. A computer node, comprising:
2 a first process residing in said node that generates a first message that grants a set
3 of tokens, if the set of tokens is available for grant, to a second process that requested
4 grant of the set of tokens, the set of tokens being required for the second process to be
5 able to modify at least one characteristic of a file stored in computer-readable memory.

1 7. A node according to claim 6, wherein:
2 each of the processes resides in a respective one of the computer nodes.

1 8. A node according to claim 7, wherein:
2 one of the processes resides in a server computer node and the other of the proc-
3 esses resides in a client computer node.

1 9. A node according to claim 6, wherein:
2 if at least one token in the set of tokens is unavailable for grant because the at
3 least one token is currently granted to a third process, the first process also generates a
4 second message that revokes current grant of the at least one token to the third process
5 prior to generating the first message.

1 10. A node according to claim 6, wherein:
2 the first message is generated by the first process in response to a request for the
3 grant of the set of tokens generated by the second process, the request specifying all to-
4 kens required for the second process to be able to modify the at least one characteristic of
5 the file.

1 11. A computer node, comprising:
2 a first process residing in said node that generates a request to a second process
3 for grant of a set of tokens required to enable the first process to modify at least one char-
4 acteristic of a file residing in computer-readable memory.

1 12. A node according to claim 11, wherein:

2 the second process resides in a second computer node, and the memory is com-
3 prised in said second node.

1 13. A node according to claim 11, wherein:
2 the set of tokens comprises all tokens required for the first process to be able to
3 modify the at least one characteristic of the file.

1 13. A network computer system, comprising:
2 a first computer node having a data file in computer-readable memory; and
3 a second computer node that issues to the first computer node a first message re-
4 questing grant of a set of tokens required to carry out a modification of at least one char-
5 acteristic of said file;
6 the first computer node issuing a second message to the second computer node
7 after receipt of the first message, the second message granting the set of tokens to the first
8 process if the set of tokens is available for grant to the second process.

1 14. A system according to claim 13, wherein:
2 the first computer node is a server node, and the second computer node is a non-
3 server node.

1 15. A system according to claim 13, wherein:
2 the set of tokens comprises all tokens required to carry out the modification of the
3 at least one characteristic of the file.

1 16. A system according to claim 13, wherein:
2 if at least one token in the set of tokens is unavailable for the grant because the at
3 least one token is currently granted, the first computer node waits to issue the first mes-
4 sage until after the first computer node receives a third message from a third computer
5 node indicating relinquishment of current grant of the at least one token.

1 17. A system according to claim 16, wherein:

2 the at least one token comprises a plurality of tokens.

1 18. Computer-readable memory containing computer-executable program instruc-
2 tions, the instructions comprising:

3 first instructions which when executed permit a data file to be maintained in com-
4 puter storage memory;

5 second instructions which when executed generate a first message requesting
6 grant of a plurality of tokens required to modify at least one characteristic of said file; and

7 third instructions which when executed generate a second message, in response to
8 said first message, that grants said tokens if said tokens are available for grant to said
9 second process.

1 19. Computer-readable memory containing computer-executable program instruc-
2 tions, the instructions comprising:

3 first instructions which when executed generate a first message that grants a set of
4 tokens, if the set of tokens is available for grant, to a requester of the set of tokens, the set
5 of tokens being required to permit the requester to be able to modify at least one charac-
6 teristic of a file stored in computer storage memory.

1 20. Computer-readable memory containing computer-executable program instruc-
2 tions, the instructions comprising:

3 first instructions that when executed generate a request for grant of a set of tokens
4 required to enable modification by an issuer of the request of at least one characteristic of
5 a file residing in storage memory.

1 21. Computer-readable memory according to Claim 18, further comprising:

2 further instructions which when executed causes, if any of said tokens are un-
3 available for grant as a result of current grant of said tokens, generation of a third mes-
4 sage revoking the current grant of said tokens.

1 22. A system according to claim 21, wherein:

2 said further instructions, in response to said third message, generate a fourth mes-
3 sage making said tokens available for grant.

1 23. Computer-readable memory according to claim 19, further comprising:
2 further instructions which when executed cause, if at least one token in the set of
3 tokens is unavailable for grant because the at least one token is currently granted, genera-
4 tion of a second message that revokes previous grant of the at least one token prior to
5 generating the first message.

1 24. Computer-readable memory according to claim 19, wherein:
2 the first message is generated in response to a request for the grant of the set of
3 tokens generated, the request specifying all tokens required to be able to modify the at
4 least one characteristic of the file.

1 25. Computer-readable memory according to claim 20, wherein:
2 the set of tokens comprises all tokens required to be able to modify the at least
3 one characteristic of the file.

1 26. A computerized data file system, comprising:
2 means for maintaining a data file in computer-readable memory; and
3 means for generating a first message requesting grant of a plurality of tokens re-
4 quired to modify at least on characteristic of said file;
5 means for generating a second message, in response to said first message, that
6 grants said tokens if said tokens are available for grant.

1 27. A system according to claim 26, further comprising:
2 means for generating, if any of said tokens are unavailable for grant as a result of
3 current grant of said tokens, a third message revoking the current grant of said tokens.

1 28. A system according to claim 27, further comprising:

2 means for generating, in response to said third message, a fourth message making
3 said tokens available for grant.

1 29. A computerized method for coherently maintaining and modifying a data file,
2 comprising:
3 maintaining a data file in computer-readable memory;
4 generating a first message requesting grant of a plurality of tokens required to
5 modify at least one characteristic of said file; and
6 generating a second message, in response to said first message, that grants said
7 tokens if said tokens are available for grant.

1 30. A method according to claim 29, further comprising:
2 if any of said tokens are unavailable for grant as a result of current grant of said
3 tokens to at least one other process, generating a third message revoking the grant of said
4 tokens.

1 31. A method according to claim 30, wherein:
2 in response to said third message, a fourth message making said tokens available
3 for grant is generated.

1 32. A computerized method for use in maintaining coherency of a data file, compris-
2 ing:
3 generating a first message that grants a set of tokens, if the set of tokens is avail-
4 able for grant, to a requester of the grant of the set of tokens, the set of tokens being re-
5 quired for requester to be able to modify at least one characteristic of the file.

1 33. A method according to claim 32, wherein:
2 if at least one token in the set of tokens is unavailable for grant because the at
3 least one token has been currently granted, the method also comprises a second message
4 that revokes current grant of the at least one token prior to generating the first message.

- 1 34. A method according to claim 32, wherein:
2 the first message is generated in response to a request for the grant of the set of
3 tokens generated by the requester, the request specifying all tokens required for the re-
4 quester to be able to modify the at least one characteristic of the file.
- 1 35. A computerized method for use in maintaining coherency of a data file, compris-
2 ing:
3 generating a request for grant of a set of tokens required to enable modification of
4 at least one characteristic of the file.
- 1 36. A method according to claim 35, wherein:
2 the set of tokens comprises all tokens required to be able to modify the at least
3 one characteristic of the file.

ABSTRACT OF THE DISCLOSURE

A computerized network file system and method are provided. The present invention permits the amount of network bandwidth that is consumed by tasks related to network file system overhead to be substantially reduced. The present invention also permits the amount of congestion in the network to be reduced.




FIG. 1
(PRIOR ART)




FIG. 2

CASE	CLIENT NODE 1	SERVER NODE	CLIENT NODE 2
1	REQUEST[A, B, C] 300	A, B, C AVAILABLE 302 GRANT [A, B, C]	NO CURRENT GRANT OF A, B, OR C
2	REQUEST[A, B] 304 C IN MEMORY (ALREADY GRANTED)	A AVAILABLE REVOKE[B] 306 310 GRANT [A, B]	B CURRENTLY GRANTED 308 RELINQUISH[B]
3	REQUEST[A, B, C] 300	REVOKE[A, B, C] 314 302 GRANT [A, B, C]	A, B, C CURRENTLY GRANTED 316 RELINQUISH[A, B, C]

FIG. 3




FIG. 4




FIG. 5




FIG. 6

DECLARATION

SOLE/JOINT INVENTOR
ORIGINAL/SUBSTITUTE/CIP

As a below named inventor, I hereby declare that: my residence, post office address, and citizenship are as stated below next to my name. I believe I am the original, first, and sole inventor (if only one name is listed below) or a joint inventor (if plural inventors are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

COMPUTERIZED FILE SYSTEM AND METHOD

as described in the specification ☒ attached or ☐ of patent Application Serial No. _____

filed _____ and amended on _____

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above; that I do not know and do not believe the same was ever known or used in the United States of America before my or our invention thereof, or patented or described in any printed publication in any country before my or our invention thereof or more than one year prior to this application; that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representative or assigns more than twelve months prior to this application; and that I acknowledge the duty to disclose information of which I am aware which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations § 1.56(a). Such information is material when it is not cumulative to information already of record or being made of record in the application, and

- (1) it establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or
- (2) it refutes, or is inconsistent with, a position the applicant has taken or may take in:

- (i) opposing an argument of unpatentability relied on by the Office, or
- (ii) asserting an argument of patentability.


I hereby claim foreign priority benefits under Title 35, United States Code § 119 of any foreign application(s) for patent or inventor's certificates listed below and have also identified below any foreign application(s) having a filing date before that of the application(s) on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE OF FILING	PRIORITY CLAIMED UNDER 35 USC 119
None			<input type="checkbox"/> YES <input type="checkbox"/> NO

I hereby claim the benefit under Title 35 United States Code § 120 of any United States application(s) listed below and, insofar as any subject matter of any claim of this application is not disclosed in the prior United States Application, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations § 1.56(a) which occurred between the filing date of the prior application and the national PCT international filing date of this application:

None

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

FULL NAME OF SOLE OR FIRST INVENTOR	INVENTOR'S SIGNATURE	DATE
Stephen Willard Dickson		10/22/99
RESIDENCE		CITIZENSHIP
31 Columbia Avenue, Nashua, NH 03060		USA
POST OFFICE ADDRESS		
Same as above		
FULL NAME OF SECOND JOINT INVENTOR	INVENTOR'S SIGNATURE	DATE
RESIDENCE		CITIZENSHIP
POST OFFICE ADDRESS		
FULL NAME OF SOLE OR FIRST INVENTOR	INVENTOR'S SIGNATURE	DATE
RESIDENCE		CITIZENSHIP
POST OFFICE ADDRESS		